Dynamic ambulance dispatching:
is the closest-idle policy always optimal?

C.J. Jagtenberg, jagtenbe@cwi.nl
S. Bhulai, s.bhulai@vu.nl
R.D. van der Mei mei@cwi.nl

May 3, 2016

Abstract

We address the problem of ambulance dispatching, in which we must decide which ambulance to send to an incident in real time. In practice, it is commonly believed that the 'closest idle ambulance' rule is near-optimal and it is used throughout most literature. In this paper, we present alternatives to the classical closest idle ambulance rule. Most ambulance providers as well as researchers focus on minimizing the fraction of arrivals later than a certain threshold time, and we show that significant improvements can be obtained by our alternative policies. The first alternative is based on a Markov decision problem (MDP), that models more than just the number of idle vehicles, while remaining computationally tractable for reasonably-sized ambulance fleets. Second, we propose a heuristic for ambulance dispatching that can handle regions with large numbers of ambulances. Our main focus is on minimizing the fraction of arrivals later than a certain threshold time, but we show that with a small adaptation our MDP can also be used to minimize the average response time. We evaluate our policies by simulating a large emergency medical services region in the Netherlands. For this region, we show that our heuristic reduces the fraction of late arrivals by 18% compared to the 'closest idle' benchmark policy. A drawback is that this heuristic increases the average response time (for this problem instance with 37%). Therefore, we do not claim that our heuristic is practically preferable over the closest-idle method. However, our result sheds new light on the popular belief that the closest idle dispatch policy is near-optimal when minimizing the fraction of late arrivals.

keywords OR in health services, Ambulances, Emergency medical services, Dispatching, Markov decision processes.

1 Introduction

Recently, Emergency Medical Services (EMS) providers have been under pressure to improve their performance. This development has led to a wide interest in ambulance planning and logistics. One issue that plays a central role is to maximize the fraction of incidents that are reached within a certain threshold time. To reach this goal without increasing the budget, solutions can be offered by operations research.

1.1 Previous Work

A large number of models are available for ambulance planning. First of all, there are models that deal with planning on a strategic level. Typically, such models determine the best locations for ambulance bases (7), and they sometimes also determine the number of vehicles that should be positioned at each base (9; 11). The majority of these solutions use mixed integer linear programming models to solve the problem. Second, there is previous work on operational ambulance planning. This seems to attract a wider range of solution methods, such as Markov chains (1), simulation-based optimization (5) and approximate dynamic programming (16).

The variety of solution methods for operational ambulance planning might be due to the difficulty of the problem. In dynamic ambulance management, the point of issue is to make decisions based on real-time information on the state of all vehicles and incidents. This makes for a complex issue, and systems quickly become intractable when the number of vehicles grows. Especially in urban areas, the situation can be considered extremely difficult because multiple vehicles operate closely to one another and therefore cannot be treated independently.

The vast majority of the papers on dynamic ambulance management have focused on how to redeploy idle vehicles, (e.g., (1; 16; 23)). Perhaps in order not to overcomplicate things, they assume a basic dispatch rule: whenever an incident occurs, they decide to send the ambulance that is closest to the incident (in time). Although this is a common dispatch policy, it was already shown to be suboptimal in 1972 (6). Regardless, most authors make this assumption without much discussion or justification; for example, the authors of (16) claim that it is an accurate enough representation of reality; however, they do not address the question of whether it is an *optimal* choice with respect to the objective (which is the fraction of incidents that is reached within the threshold time). The authors of (1) do not address the assumption at all.

The 'closest idle ambulance' dispatch policy appears to be so natural that justification is unnecessary, but one should not overlook the possibility to change the dispatch policy in order to improve the objective. The authors of (23) recognize that this is a possibility. Nevertheless, they focus on relocating the vehicles when they become idle, instead of when they are dispatched. It should be clear that an EMS system can benefit from an improved dispatch policy, but since the topic has been underexposed in current literature, it is still unknown how much benefit can be expected. Furthermore, a dispatch policy can be combined with a relocation rule to realize even larger improvements in the objective value.

Few papers have discussed dispatch rules other than sending the closest idle ambulance. One exception is the paper of (5), in which the authors divide the region into separate sub-regions, and each sub-region has a list of stations from which a vehicle should preferably depart. Another example is (21), which compares two different dispatch rules; the so-called 'closest-ambulance response' versus 'regionalized response'. Under regionalized response, each ambulance serves its own region first, even if it is temporarily outside its region. Only if it is unavailable, the closest idle ambulance is sent. However, note that both examples still ignore important information: the outcome does not depend on whether some regions remain uncovered after the dispatch is performed. Alternatively, a choice could be made such that the remaining idle vehicles are in a good position with respect to expected incidents in the near future. This ensures that future incidents get a larger likelihood of being reached in time, thereby increasing the total expected fraction of incidents that can be reached within the time threshold.

There are a few authors who have used MDPs to solve the dispatch problem. In (13) the authors define costs in their MDP, but they do not discuss the meaning or interpretation of this. In their numerical work, they use randomly drawn instances – and although it is not mentioned explicitly, it appears also these costs are drawn at random. Furthermore, they do not compare their solution with the closest-idle method. The authors of (2) maximize patient survivability. We conclude that neither (13) nor (2) analyzes the fraction of late arrivals.

Building on (2), there is a series of papers that considers a dispatch problem with prioritized patients (17; 18). The main idea is to allow increased response times for the non-urgent patients, such that shorter response times can be realized for the urgent patients. Although this approach makes sense from a practical point of view, it is not the goal of our paper. Instead, we assume that *all* patients have high priority, and ask the question how to dispatch vehicles such that the fraction of late arrivals is maximized. Furthermore, (17) and (18) do not discuss the scalability of their approach, and include numerical work for just 4 vehicles and 4 demand locations. Note that this is significantly smaller than the regions we have in mind: the Netherlands is divided into 24 EMS regions, varying in size between 40 and 456 postal codes.

One paper explicitly opposes the idea of dispatching an ambulance other than the closest idle one (10). However, the authors do not show that the closest idle method is optimal; in fact, they only show that the dispatch rule used in practice performs worse than the 'closest idle ambulance rule' would. Based on their description, the reason for the poor performance in practice seems to be a computer-aided dispatch system that is not accurate enough to determine the true positions of the vehicles. We emphasize that accurate location information is crucial in order to determine the best ambulance to send to an incident. Throughout this paper, we will assume that such information is present. In many regions, such as Utrecht in the Netherlands, a monitoring tool is available that refreshes the exact GPS coordinates of each vehicle every 30 seconds. This seems accurate enough for our purposes.

1.2 Our Contribution

The main goal of this paper is to better understand the ambulance dispatch process. In particular, we question the often-made assumption that one cannot do much better than the 'closest idle' dispatch method.

Thereto, we search for sensible dispatch rules other than the classical closest idle ambulance policy. We present different dispatch rules that outperform the closest idle policy for different objectives.

First, we propose a Markov Decision Problem (MDP) for ambulance dispatching, where the state space is described by an optional incident location and the availability of the ambulances. We mainly focus on minimizing the fraction of arrivals later than a target time – a typical objective in ambulance planning. However, we show that with a small change, our model can also minimize the average response time.

Second, we propose a heuristic for ambulance dispatching that behaves similar to the policy obtained from the MDP. However, it is able to determine more accurately what the response time would be when dispatching a driving ambulance. Furthermore, the heuristic can be computed in polynomial time, which allows us to apply it to regions with a large number of vehicles.

We validate our policies by a discrete-event simulation model of an urban EMS region. These simulations indicate that our proposed dispatch heuristic can decrease the fraction of late arrivals by as much as 18% relatively compared to the closest idle ambulance dispatch method. Our result sheds new light on the popular belief that deviating from the closest idle policy cannot greatly improve the objective. In the field of ambulance management, an improvement of 18% is considered large; however, it should be noted that there is a tradeoff: our policy significantly increases the average response time. Although we do not advise all EMS managers to immediately discard the closest idle dispatch method, we do show that the typical argument – that it would not lead to large improvements in the fraction of late arrivals – should be changed.

The rest of this paper is structured as follows. In Section 2, we give a formal problem definition. In Section 3, we present our proposed solution using Markov Decision Processes (MDPs), followed by a solution based on a scalable heuristic in Section 4. We show our results for a small, intuitive region in Section 6 and in a realistic case study for the area of Utrecht in Section 7 and Section 8. We end with our conclusions in Section 9.

2 Problem Formulation

Define the set V as the set of locations at which incidents can occur. Note that these demand locations are modeled as a set of discrete points. Incidents at locations in V occur according to a Poisson process with rate λ.[1] Let d_i be the fraction of the demand rate λ that occurs at node i, $i \in V$. Then, on a smaller scale, incidents occur at node i with rate λd_i.

Let A be the set of ambulances, and $A_{idle} \subseteq A$ the set of currently idle ambulances. When an incident has occurred, we require an idle ambulance to immediately drive to the scene of the incident. The decision which ambulance to send has to be made at the moment we learn about the incident, and is the main question of interest in this paper. When an incident occurs and there are no idle ambulances, the call goes to a first-come first-serve queue. Note that when an incident occurs and an ambulance is available, it is not allowed to postpone the dispatch. Although in some practical situations dispatchers may queue a low priority call when the number of idle servers is small, in this paper we will focus on the most urgent incidents, which require service immediately.

Our objectives are formulated in terms of response times: the time between an incident and the arrival of an ambulance. In practice, incidents have the requirement that an ambulance must be present within T time units. Therefore, we want to minimize the fraction of incidents for which the response time is larger than T. Another observation is that we want response times to be short, regardless of whether they are smaller or greater than T. We translate this into a separate objective, which is to minimize the average response time. We assume that the travel time $\tau_{i,j}$ between two nodes $i, j \in V$ is deterministic, and known in advance.

Sending an ambulance to an incident is followed by a chain of events, most of which are random. When an ambulance arrives at the incident scene, it provides service for a certain random time τ_{on_scene}. Then it is decided whether the patient needs transport to a hospital. If not, the ambulance immediately becomes idle. Otherwise, the ambulance drives to the nearest hospital in the set $H \subseteq V$. Upon arrival, the patient is transferred to the emergency department, taking a random time $\tau_{hospital}$, after which the ambulance becomes idle.

An ambulance that becomes idle may be dispatched to another incident immediately. Alternatively, it may return to its base location. Throughout this paper, we will assume that we are dealing with a *static*

[1] We will discretize the arrival process in the next section.

V	The set of demand locations.
H	The set of hospital locations, $H \subseteq V$.
A	The set of ambulances.
A_{idle}	The set of idle ambulances.
W_a	The base location for ambulance a, $a \in A$, $W_a \in V$.
T	The time threshold.
λ	incident rate.
d_i	The fraction of demand in i, $i \in V$.
$\tau_{i,j}$	The driving time between i and j with siren turned on, $i, j \in V$.

Table 1: Notation.

ambulance system, i.e., each ambulance has a fixed, given base and may not drive to a different base. However, it is possible that multiple ambulances have the same base location. We denote the base location of ambulance a by W_a, for $a \in A$.

An overview of the notation can be found in Table 1.

3 Solution Method: Markov Decision Process

We model the ambulance dispatch problem as a discrete-time Markov Decision Process (MDP). In each state s (further defined in Section 3.1), we must choose an action from the set of allowed actions: $\mathcal{A}_s \subseteq \mathcal{A}$, which we describe in Section 3.2. The process evolves in time according to transition probabilities that depend on the chosen actions, as described in Section 3.4. We are dealing with an infinite planning horizon, and our goal is to maximize the average reward. We eventually find our solution by performing value iteration (20). Our choice to use value iteration was motivated by it being simple in implementation, and sufficient to answer our central question on the closest idle policy.

In our model, we assume that at most one incident occurs within a time step. Therefore, the smaller the time steps, the more accurate the model will be. However, there is a tradeoff, as small time steps will increase the computation time. Throughout this paper, we take time steps to be one minute, which balances the accuracy and the computation time.

3.1 State Space

When designing a state space, it is important to store the most crucial information from the system in the states. However, when dealing with complex problems – such as real-time ambulance planning – it is tempting to store so much information, that the state space becomes intractable. This would lead to the so-called curse of dimensionality (3), which makes it impossible to solve the problem with well-known Markov Decision Problem (MDP) approaches.

As discussed before, there is little previous work on how to choose a good dispatch policy, but to some extent we can draw parallels with work on dynamic ambulance redeployment (which relocates idle vehicles): some researchers overcome the problem of an intractable state space by turning to Approximate Dynamic Programming, which allows for an elaborate state space to be solved approximately (16). Alternatively, some researchers choose a rather limited state space, for example, by describing a state merely by the *number* of idle vehicles (1).

For our purpose, i.e., to determine *which* ambulance to send, it is important to know whether the ambulance we might send will arrive within T time units. Therefore, it is crucial to know where the incident took place. Furthermore, we require some knowledge of where the idle ambulances are. Clearly, storing only the number of idle vehicles would be insufficient. However, storing the location of each idle ambulance would already lead to an intractable state space for practical purposes. Instead, we can benefit from the fact that we are trying to improve a *static* solution. In a static solution, the home base for any ambulance is known in advance. Note that an idle ambulance must be either residing at its base location, or travelling towards the base. Hence, if we allow for an inaccuracy in the location of idle ambulances, in the sense that we use their destination rather than their actual location, their location does not need to be part of the state. Merely keeping track of a simple status for each ambulance (idle or not), now suffices. Thereto, let $stat_i$ denote this status for ambulance i:

$$stat_i \in \{idle, busy\}, \quad \forall i \in A.$$

This leads us to a state s, defined as follows.

$$(Loc_{acc}, stat_1, stat_2, \ldots, stat_{|A|}), \tag{1}$$

where Loc_{acc} denotes the location of the incident that has just occurred in the last time step. In case no incident occurred in the last time step, we denote this by a dummy location, hence

$$Loc_{acc} \in V \cup \{0\}.$$

This leads to a state space of size $(|V| + 1)2^{|A|}$.

For future reference, let $Loc_{acc}(s)$ denote the location of the incidents that have occurred in the previous time step when the system is in state s. For ease of notation, we introduce boolean variables $idle_i(s)$ and $busy_i(s)$ to denote whether $stat_i$ is idle or busy in state s, $\forall i \in A, \forall s \in S$.

3.2 Policy Definition

In general, a policy Π can be defined as a mapping from the set of states to a set of actions: $S \to \mathcal{A}$. In our specific case, we define $\mathcal{A} = A \cup \{0\}$; that is if $\Pi(s) = a$, for $a \in A$, ambulance a should be sent to the incident that has just occurred at $Loc_{acc}(s)$. Action 0 may be interpreted as sending no ambulance at all (this is typically the choice when no incident occurred in the last time step, or when no ambulance is available).

In a certain state, not all actions are necessarily allowed. Denote the set of feasible actions in state s as

$$\mathcal{A}_s \subseteq \mathcal{A}, \quad \forall s \in S.$$

For example, it is not possible to send an ambulance that is already busy with another incident. This implies

$$busy_a(s) \to a \notin \mathcal{A}_s, \quad \forall a \in A, \quad \forall s \in S. \tag{2}$$

Furthermore, let us require that when an incident has taken place, we must always send an ambulance – if any are idle.

$$\exists a \in A : idle_a(s) \ \wedge \ Loc_{acc}(s) \neq 0 \to 0 \notin \mathcal{A}_s, \quad \forall s \in S. \tag{3}$$

Moreover, if no incident has occurred, we may simplify our MDP by requiring that we do not send an ambulance:

$$Loc_{acc}(s) = 0 \to \mathcal{A}_s = \{0\}, \quad \forall s \in S. \tag{4}$$

All other actions from $\mathcal{A}$ that are not restricted by (2)–(4) are feasible. This completely defines the allowed action space for each state.

3.3 Rewards

In ambulance planning practice, a typical goal is to minimize the fraction of late arrivals. Since our decisions have no influence on the number of incidents, this is equivalent to minimizing the *number* of late arrivals. An alternative goal might be to minimize average response times. Our MDP approach may serve either of these objectives, simply by changing the reward function.

Define $R(s, a)$ as the reward received when choosing action a in state s, $\forall s \in S, \forall a \in \mathcal{A}_s$. Note that in this definition, the reward does not depend on the next state. Keep in mind that our goal is to maximize the average rewards.

3.3.1 Fraction of Late Arrivals

To minimize the fraction of late arrivals, i.e., the fraction of incidents for which the response time is greater than T, we define the following rewards:

$$R(s,a) = \begin{cases} 0 & \text{if } Loc_{acc}(s) = 0; \\ -N & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a = 0, \\ & \text{i.e., no idle ambulances;} \\ 0 & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a \in A \\ & \quad \wedge \ \tau_{W_a,\ Loc_{acc}(s)} \leq T; \\ -1 & \text{otherwise.} \end{cases}$$

Here N is a number that is typically greater than 1. We discuss the choice of this parameter further in Section 3.6

3.3.2 Average Response Time

To minimize the average response time, one may use the same MDP model, except with a different reward function. Let M be a large enough number, typically such that $M > \tau_{i,j}, \forall i, j \in V$. Then we can define the rewards as follows.

$$R(s,a) = \begin{cases} 0 & \text{if } Loc_{acc}(s) = 0; \\ -M & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a = 0,; \\ & \text{i.e., no idle ambulances;} \\ -\tau_{W_a,\ Loc_{acc}(s)} & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a \in A. \end{cases}$$

3.4 Transition Probabilities

Denote the probability of moving from state s to s', given that action a was chosen, as:

$$p^a(s,s'), \qquad \forall a \in \mathcal{A}_s, \qquad \forall s, s' \in S.$$

To compute the transition probabilities, note that the location of the next incident is independent of the set of idle ambulances. Thereto, $p^a(s,s')$ can be defined as a product of two probabilities. We write $p^a(s,s') = P_1(s') \cdot P_2^a(s,s')$, which stands for the probability that an incident happened at a specific location (P_1), and the probability that specific ambulances became available (P_2), respectively.

First of all, let us define $P_1(s')$. Since incidents occur according to a Poisson process, we can use the arrival rate λ (the probability of an arrival anywhere in the region per discrete time step) to obtain

$$P_1(s') = \begin{cases} \lambda \cdot d_{Loc_{acc}(s')} & \text{if } Loc_{acc}(s') \in V; \\ 1 - \lambda & \text{else.} \end{cases}$$

Note that the occurrence of incidents does not depend on the previous state (s).

Secondly, we need to model the process of ambulances that become busy or idle. For tractability, we will define our transition probabilities as if ambulances become idle according to a geometric distribution. In reality – and in our verification of the model – this is not the case, but since our objective is the long term average cost, this modelling choice leads to the same performance. Let us define a parameter $r \in [0,1]$, which represents the rate at which an ambulance becomes idle. We discuss the parameter choice for r in Section 3.6.

We include a special definition if an ambulance was just dispatched. In such a case, the ambulance cannot be idle in the next time step. Furthermore, ambulances do not become busy, unless they have just been dispatched.

We now define

$$P_2^a(s,s') = \Pi_{i=1}^{|A|} P_{change}^a\big(stat_i(s), stat_i(s')\big), \quad \forall s, s' \in S,$$

where

$$P^a_{change}\big(stat_i(s), stat_i(s')\big) = \begin{cases} 1 & \text{if } a = i \ \wedge \ busy_i(s'); \\ 0 & \text{if } a = i \ \wedge \ idle_i(s'); \\ r & \text{if } a \neq i \ \wedge \ busy_i(s) \\ & \wedge \ idle_i(s'); \\ 1-r & \text{if } a \neq i \ \wedge \ busy_i(s) \\ & \wedge \ busy_i(s'); \\ 0 & \text{if } a \neq i \ \wedge \ idle_i(s) \\ & \wedge \ busy_i(s'); \\ 1 & a \neq i \ \wedge \ idle_i(s) \\ & \wedge \ idle_i(s'). \end{cases} \tag{5}$$

3.5 Value Iteration

Now that we have defined the states, actions, rewards and transition probabilities, we can perform value iteration to solve the MDP. Value iteration, also known as backward induction, calculates a value $V(s)$ for each state $s \in S$. The optimal policy, i.e., the best action to take in each state, is the action that maximizes the expected value of the resulting state s'.

$V(s)$ is calculated iteratively, starting with an arbitrary value $V_0(s)\ \forall s \in S$. (In our case, we start with $V_0(s) = 0\ \forall s \in S$.) In each iteration i, one computes the values $V_i(s)$ given $V_{i-1}(s)\ \forall s \in S$ as follows.

$$V_i(s) := \max_{a \in \mathcal{A}_s} \{ \sum_{s'} p^a(s, s')(R(s, a) + V_{i-1}(s')) \} \tag{6}$$

This is known as the 'Bellman equation' (4).

When the span of V_i, i.e., $\max V_i(s) - \min V_i(s)$, converges, the left-hand side becomes equal to the right-hand side in Equation (6), except for an additive constant. After this convergence is reached, the value of $V(s)$ is equal to $V_i(s)\ \forall s \in S$. Note that the MDP we defined is unichain. Hence, value iteration is guaranteed to converge.

Small regions, such as the region in Section 6, allow us to reach convergence and accurately determine the value function V. However, for larger regions (such as Utrecht in Section 7), value iteration simply takes too much time to reach convergence. Instead, we use the non-converged values V_i and analyze the performance of the corresponding policy.

3.6 Parameter choices

Recall that $-N$ is the reward given in the situation that there occurs an incident while all ambulances are busy, in the MDP that attempts to minimize the fraction of late arrivals. If $N > 1$, this implies that when all ambulances are busy, the rewards are smaller than when we send an ambulance that takes longer than T to arrive. This is in agreement with the general idea that having no ambulances available is a very bad situation. One might be tempted to make the reward for the only possible action ($a = 0$) in these states even smaller than we did, in order to influence the optimal actions in other states: the purpose would be to steer the process away from states with no ambulances available. However, note that this would not be useful, because our actions do not affect how often we end up in a state where all ambulances are busy. This is merely determined by the outcome of an external process, i.e., an unfortunate sequence of incidents. Therefore, an extremely small reward for action $a = 0$ in states where all ambulances are busy, would only blur the differences between rewards for actions in other states. (In our numerical experiments, we use $N = 5$.)

For the MDP that minimizes the average response times, the reward given in the situation that there occurs an incident while all ambulances are busy is given by $-M$. In our numerical experiments, we use $M = 15$ for the small region, and $M = 30$ for the region Utrecht. (In our implementation, time steps are equal to minutes.)

Recall that r is the rate at which an ambulance becomes idle. We should set it in such a way, that the expected duration is equal to the average in practice. So this includes an average travel time, and an average time spent on scene. We add an average driving time to a hospital to that, as well as a realistic hospital drop off time – both multiplied with the probability that a patient needs to go to the hospital. For Dutch ambulances, this results in an average of roughly 38 minutes to become available after departing to an incident. For the geometric distribution, we know that the maximum likelihood estimate $\hat{r}$ is given by one divided by the sample mean. In this case, $\hat{r} = \frac{1}{38} \approx 0.0263$, which we use as the value for r in our numerical experiments.

4 Solution Method: dynamic MEXCLP heuristic for dispatching

In this section we describe a dispatch heuristic that is easy to implement and scales well. It can be computed in real time, for any number of vehicles and ambulance bases that is likely to occur in practice. The method is inspired by dynamic MEXCLP, a.k.a. 'DMEXCLP' (12), a heuristic for real-time redeployment of ambulances.

The general idea is that, at any time, we can calculate the *coverage* provided by the currently idle ambulances. This results in a number that indicates how well we can serve the incidents that might occur in the (near) future.

More specifically, coverage is defined as in the MEXCLP model (9), that we will describe next.

4.1 Coverage according to the MEXCLP model

In this section we briefly describe the objective of the well-known MEXCLP model. MEXCLP was originally designed to optimize the distribution of a limited number, say $|A|$, ambulances over a set of possible base locations W. Each ambulance is modeled to be unavailable with a pre-determined probability q, called the *busy fraction*. Consider a node $i \in V$ that is within range of k ambulances. The travel times $\tau_{i,j}$ $(i, j \in V)$ are assumed to be deterministic, which allow us to straightforwardly determine this number k. If we let d_i be the demand at node i, the expected covered demand of this vertex is $E_k = d_i(1-q^k)$. Note that the marginal contribution of the kth ambulance to this expected value is $E_k - E_{k-1} = d_i(1-q)q^{k-1}$. Furthermore, the model uses binary variables y_{ik} that are equal to 1 if and only if vertex $i \in V$ is within range of at least k ambulances. The objective of the MEXCLP model can now be written as:

$$\text{Maximize} \sum_{i \in V} \sum_{k=1}^{|A|} d_i(1-q)q^{k-1} y_{ik}.$$

In (9), the author adds several constraints to ensure that the variables y_{ik} are set in a feasible manner. For our purpose, we do not need these constraints, as we shall determine how many ambulances are within reach of our demand points – the equivalent of y_{ik} – in a different way.

4.2 Applying MEXCLP to the dispatch process

The dispatch problem requires us to decide which (idle) ambulance to send, at the moment an incident occurs. Thereto, we compute the *marginal* coverage that each ambulance provides for the region. The ambulance that provides the smallest marginal coverage, is the best choice for dispatch, in terms of remaining coverage for future incidents. However, this does not incorporate the desire to reach the current incident within target time T. We propose to combine the two objectives – reaching the incident in time and remaining a well-covered region – by always sending an ambulance that will reach the incident in time, if possible. This still leaves a certain amount of freedom in determining which *particular* ambulance to send.

The computations require information about the location of the (idle) ambulances. Denote this by $Loc(a)$ for all $a \in A_{idle}$. We evaluate two different options for $Loc(a)$, that we describe next.

Using real positions of ambulances is the most accurate information one could use. In practice, $Loc(a)$ may be determined by GPS signals. For simulation purposes, the current position of the ambulance while driving may be determined using, e.g., interpolation between the origin and destination, taking into account the travel speed. In either case, the location should be rounded to the nearest point in V, because travel times $\tau_{i,j}$ are only known between any $i, j \in V$.

Using destinations of ambulances is a far simpler, albeit somewhat inaccurate alternative. The simplicity, however, does make it a practical and accessible option. When determining $Loc(a)$, simply take the destination of ambulance a. This is a good option, e.g., when no – or not enough – GPS information is available. Furthermore, this solution has a certain fairness in comparison to the MDP solution in Section 3, which is also required to make decisions based on the destinations of ambulances.

Let A^+_{idle} denote the set of idle ambulances that are able to reach the incident in time, i.e., the ambulances $a \in A_{idle}$ for which $\tau_{Loc(a),i} \leq T$ (where i denotes the incident location). Note that this definition depends on how $Loc(a)$ was chosen: when based on the true locations of ambulances, the set A^+_{idle} can be determined correctly. When one uses the destinations of ambulances, the decision of which ambulances are in A^+_{idle} may contain errors: some ambulances may in fact be closer to the incident than they appear (because they are driving towards a base that is further away from the incident), or the other way around they may in reality be further away from the incident than Loc_a suggests.

Similarly, let A^-_{idle} denote the set of idle ambulances that cannot reach the incident in time, which implies that $A^+_{idle} \cup A^-_{idle} = A_{idle}$. Then, if $A^+_{idle} \neq \emptyset$, we decide to dispatch a vehicle that will arrive within the threshold time, but chosen such that the coverage provided by the remaining idle vehicles is as large as possible:

$$\arg\min_{x \in A^+_{idle}} \sum_{i \in V} d_i (1-q) q^{k(i, A_{idle})-1} \cdot \mathbb{1}_{\tau_{Loc(x),i} \leq T}. \tag{7}$$

Otherwise, simply dispatch a vehicle such that the coverage provided by the remaining idle vehicles is as large as possible (without requiring an arrival within the threshold time):

$$\arg\min_{x \in A^-_{idle}} \sum_{i \in V} d_i (1-q) q^{k(i, A_{idle})-1} \cdot \mathbb{1}_{\tau_{Loc(x),i} \leq T}. \tag{8}$$

Note that in our notation, k is a function of i and A_{idle}. $k(i, A_{idle})$ represents the number of idle ambulances that are currently within reach of vertex i. After choosing the locations of ambulances that one wishes to use – the real locations or the destinations – $k(i, A_{idle})$ can be counted in a straightforward manner.

We have seen that the way one measures the location of ambulances – either the true location or just the destination – affects the definition of the set A^+_{idle} (resp. A^-_{idle}), and thereby also the number $k(i, A_{idle})$ in Equation 8. There is, however, one more aspect that is affected by the location of the ambulance: this is incorporated in $\mathbb{1}_{\tau_{Loc(x),i} \leq T}$ in Equation 8. Hence, using the destination of ambulances results in a small error in three different places. It is reasonable to assume that using the destinations of ambulances performs worse than using the real locations, but the magnitude of the performance difference is hard to oversee beforehand. Instead, we will show the performance difference in retrospect in our numerical examples in Sections 6, 7 and 8.

5 Simulation model

In order to compare the results of different policies, we measure their performance using simulation. All results mentioned in this paper, including the fraction of late arrivals and the average response times, are estimates based on the observed response times in our simulation model. This section describes the details and assumptions of the simulation model, and in particular highlight how they differ from the MDP.

The reason for using simulation is that the EMS process is rather complex. The aforementioned MDP does not capture all details and is therefore not able to estimate the performance accurately. We will next describe the two main differences between the MDP and the simulation.

One reason why the MDP is not entirely accurate, is that incidents that occur while no vehicles are available are 'lost'. This assumption is made to improve scalability: it avoids the need to expand the state with a queue of calls that are waiting for vehicles to become idle. However, counting these calls as lost is technically incorrect for two reasons: first of all, an ambulance might become available shortly after, and it is – although unlikely – still possible that it arrives within the time threshold. Second, a lost call in the MDP is not counted in the total workload, which leads to an overestimation in the number of idle vehicles in the time steps shortly after the lost call. In our simulation, we place the incidents that arrive while all vehicles are busy in a first come first serve queue. Ambulances that become idle are immediately dispatched to a waiting incident (if any), or else head back to their home base.

Our simulations are also able to handle the rather complex course of events that take place when an ambulance is dispatched while on the road. Such vehicles are typically returning to the home base, already idle and ready to serve an incoming call. Our simulation computes the current location of the vehicle based on an interpolation between the origin (typically a hospital where a patient was just dropped off) and the destination (the vehicle's home base) of the trip, taking into account the total time of that particular trip. The MDP is unable to distinguish between idle vehicles on the road and vehicles at the base. Adding on-the-road information to the MDP would require a state definition that includes (at least) the drop off location of the last patient. This alone would already lead to a state space explosion and therefore we do not recommend solving this for realistic instances.

In our simulation, $\tau_{onscene}$ is exponentially distributed with an expectation of 12 minutes. $\tau_{hospital}$ is drawn from a Weibull distribution with an expectation of 15 minutes. More specifically, it has shape parameter 1.5 and scale parameter 18 (in minutes). We state these distributions for completeness, however, numerical experiments (done by the authors in ongoing work) indicate that the performance of dynamic ambulance management does not depend much on the chosen distribution for $\tau_{onscene}$ or $\tau_{hospital}$, and we conjecture that the same holds for the dispatching problem. In our simulations, patients need hospital treatment with probability 0.8. This value was estimated from Dutch data (8). (Similar numbers (78% nation-wide) can be deduced from (19).)

Note that $\tau_{onscene}$ or $\tau_{hospital}$ and the probability that a patient needs hospital treatment are not explicitly part of our solution methods. Instead, they subtly affect the busy fraction q (for the heuristic) or the transition probabilities with rate r (for the MDP).

6 Results: A motivating example

In this section, we consider a small region for which there is some intuition with respect to the best dispatch policy. We show that the intuitive dispatch policy that minimizes the fraction of late arrivals, is in fact obtained by both our solution methods (based on MDP and MEXCLP). We will address the alternative objective, i.e., minimizing the average response times, as well.

Figure 1 shows a toy example for demonstrative purposes. We let calls arrive according to a Poisson process with on average one incident per 45 minutes. Furthermore, incidents occur w.p. 0.1 in Town 1, and w.p. 0.9 in Town 2. Eighty percent of all incidents require transport to the hospital, which is located in Town 2.

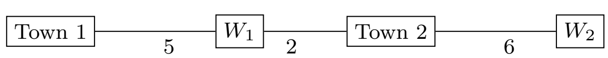

Figure 1: A graph representation of the region. The numbers on the edges represent the driving times in minutes with siren turned on. W_1 and W_2 represent the base locations of ambulance 1 and 2, respectively. Incidents occur only in Town 1 and Town 2. The only hospital is located in Town 2.

6.1 Fraction of late arrivals

This section deals with minimizing the fraction of response times greater than 12 minutes. A quick analysis of the region in Figure 1 leads to the observation that the 'closest idle' dispatch strategy must be suboptimal. In order to serve as many incidents as possible within 12 minutes, it is evident that the optimal dispatch strategy should be as follows: when an incident occurs in Town 2, send ambulance 2 (if available). When an incident occurs in Town 1, send ambulance 1 (if available). Both the MDP solution that attempts to minimize the fraction of late arrivals (with, e.g., $N = 5$), as well as the dispatch heuristic based on MEXCLP, lead to this policy.

The response times obtained by simulating the closest-idle policy and MDP (frac) solution are compared in Figure 2. This clearly shows that the MDP solution outperforms the closest idle method, as was expected.

Note that in our model, it is mandatory to send an ambulance, if at least one is idle. Furthermore, our proposed solutions do not base their decision on the locations of idle ambulances (instead, we pretend they are at their destination, which is fixed for each ambulance). Therefore, in this example with 2 ambulances, one can describe a dispatch policy completely by defining which ambulance to send when both are idle, for each possible incident location. For an overview of the various policies, see Table 2. As shown in this table, the MDP solution minimizing the fraction of late arrivals – in this particular instance – comes down to exactly the same policy as the MEXCLP dispatch heuristic using destinations of vehicles. Therefore, the results mentioned for either of those two policies, also hold for the other. For this problem instance the closest-idle dispatch method turns out to be roughly equivalent with the MDP solution minimizing the average response time (except for the fact that the MDP can only use destinations of vehicles, whereas closest-idle uses their true positions).

solution method	$Loc_{acc} = Town1$	$Loc_{acc} = Town2$
MEXCLP(dest)	W_1	W_2
MDP(frac)	W_1	W_2
MDP(avg)	W_1	W_1

Table 2: An overview of the behaviour of various dispatch policies when both ambulances are idle. The value in the table represents the base from which an ambulance should be dispatched.

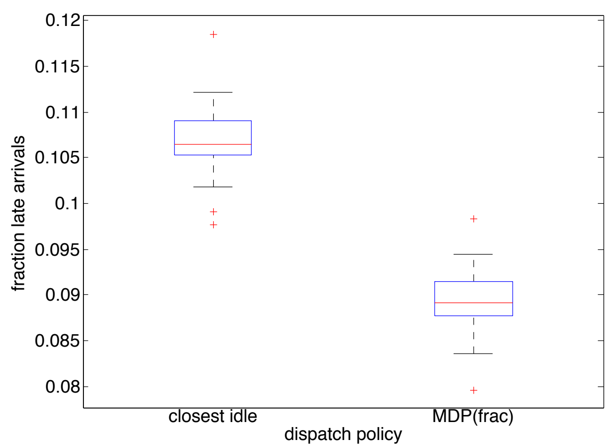

Figure 2: Box plots showing the fraction of late arrivals as observed in a simulation of the small region. This figure shows the performance of the MDP solution that attempts to minimize the fraction of late arrivals (after value iteration converged). The performance is compared with the 'closest idle' dispatch policy. Each policy was evaluated with 20 runs of 5,000 simulated hours. The red plus-signs indicate outliers.

6.2 Average response time

We used the MDP method described in Section 3.3.2 to obtain a policy that should minimize the average response time, let us denote this policy by MDP(avg). We evaluate the performance of the obtained policy, again by simulating the EMS activities in the region. These simulations show that the MDP solution indeed reduces the average response time significantly, compared to the policy that minimizes the fraction of late arrivals (MDP(frac)) – see Figure 3.

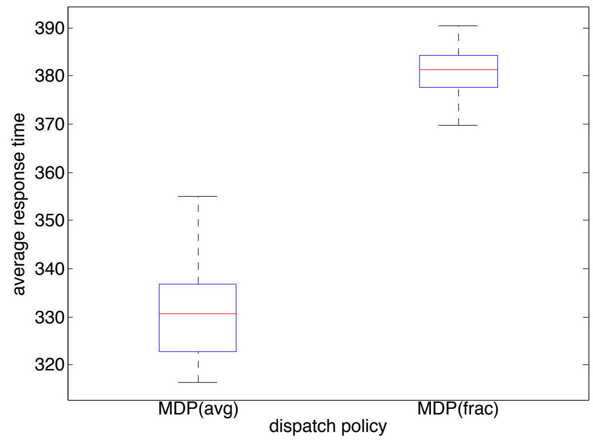

Figure 3: Box plots showing the average response times in seconds, as observed in simulations of the small region. This figure shows the performance of the MDP solution that attempts to minimize the fraction of late arrivals versus the MDP solution that attempts to minimize the average response time (after value iteration has converged). Each policy was evaluated with 20 runs of 5,000 simulated hours. The red plus-signs indicate outliers.

7 Results: Region Utrecht with 8 vehicles

In this section, we validate our redeployment method on a realistic problem instance.

7.1 The region

As in (12), we modeled the region of Utrecht, which is hosted by one of the largest ambulance providers of the Netherlands. Note that this is not one city, but the *county* of Utrecht: an area which contains several cities, including Amersfoort and Utrecht city. However, the whole region may – by international standards – be considered an urban area. For the parameters used in the implementation, see Table 3. This is a region with multiple hospitals, and for simplicity we assume that the patient is always transported to the nearest hospital, if necessary.

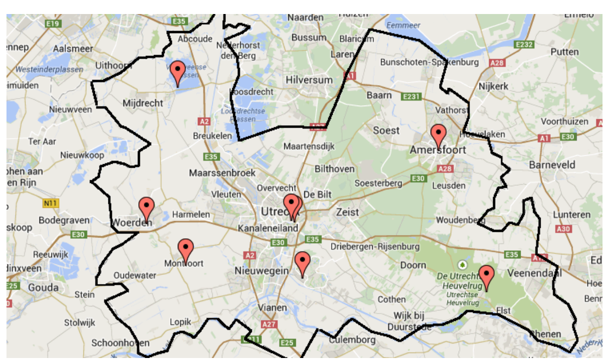

Figure 4: The home bases for each of the 8 ambulances in region Utrecht. The chosen locations currently exist as base locations operated by the ambulance provider for this region. Note that in this figure, two vehicles are stationed at the base in the center of Utrecht.

Note that we used the fraction of inhabitants as our choice for d_i. In reality, the fraction of demand could differ from the fraction of inhabitants. However, the number of inhabitants is known with great accuracy, and this is a straightforward way to obtain a realistic setting. Furthermore, the analysis of robust optimization

parameter	**magnitude**	**choice**
A	8	Small enough for a tractable MDP.
λ	1/15 minutes	A reasonable number of incidents for 8 ambulances.
W_a $(a \in A)$		Postal codes 3582, 3645, 3958, 3582, 3991, 3447, 3811, 3417.
V	217	4 digit postal codes.
H	10	The hospitals within the region in 2013, excluding private clinics.
$\tau_{i,j}$		Driving times as estimated by the RIVM.
d_i		Fraction of inhabitants as known in 2009.

Table 3: Parameter choices for our implementation of the region of Utrecht.

for uncertain ambulance demand in (14) indicates that we are likely to find good solutions, even if we make mistakes in our estimates for d_i.

In the Netherlands, the time target for the highest priority emergency calls is 15 minutes. Usually, 3 minutes are reserved for answering the call, therefore we choose to run our simulations with $T = 12$ minutes. The driving times for EMS vehicles between any two nodes in V were estimated by the Dutch National Institute for Public Health and the Environment (RIVM) in 2009. The RIVM uses measurements of a full year of ambulance movements for this, and differentiates between road type, region and time of day. The driving times we use are estimates we for ambulance movements with the siren turned on, at the time of day with most traffic congestion. Therefore, they could be considered a pessimistic or safe approximation. Note that these travel times are deterministic. For ambulance movements without siren, e.g., when repositioning, we used 0.9 times the speed with siren. The locations we used as home bases are depicted in Figure 4, and correspond to actual base locations in the EMS region.

The number of vehicles used in our implementation is such that value iteration is still tractable. For this problem instance, one iteration takes approximately 70 minutes to calculate. Although this seems to be rather long, we emphasize these calculations take place in a preparatory phase. We perform 21 iterations after which the current solution is used as policy. After these calculations, the final policy constitutes a lookup table for which online decision making can be done without additional computation time.

7.2 Analysis of the MDP solution

In this section, we highlight and explain some features of the MDP solution that attempts to minimize the fraction of late arrivals for the region Utrecht. In particular, we will focus on the states for which the MDP solution differs from the closest idle policy.

The output of the MDP is a table with the incident location, the status of the different ambulances (idle or not), and the optimal action. This output is a rather large table that does not easily reveal insight into the optimal policy. Therefore, we used classification and regression trees (CART trees) on the table to find structure in the form of a decision tree. We used random forests to create the decision tree, since it is known that a basic CART has poor predictive performance (see Chapter 14 of (15)). Another option is to used bagging (i.e., bootstrap aggregation) trees. This effectively generates several bootstrap samples of the original table, trains CART trees on the sample, and finally averages the results. While bagging trees reduces the variance in the prediction, random forests also cancel any correlation structure in the generation of the trees that may present while bagging.

The outcome that describes the best policy after 21 value iterations is a decision tree that divides the

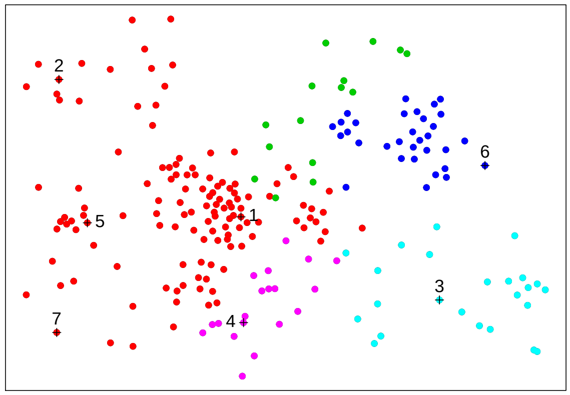

Figure 5: Each node represents a postal code in Utrecht. Nodes with the same colour have similar MDP solutions. The numbers indicate the ambulance bases. (Two vehicles are stationed at base number 1.)

state space into five regions, see Figure 5. If an incident occurs in the red region, then in most cases the closest idle ambulance is dispatched. If the ambulance at base 4 is idle, this is even more often the case than when it is busy. The location of the base stations plays an essential role in the final decision tree.

For some nodes, whether or not the closest idle ambulance should be dispatched depends even more heavily on which ambulances are idle. For example, if an incident occurs in the dark blue region while the ambulance at base 6 is idle, the MDP tells us to almost always (in more than 98% of the states) send the closest idle ambulance. Conversely, if the ambulance at base 6 is busy, it is better to strategically choose a different ambulance instead of simply applying the closest idle policy.

This may be intuitively understood as follows. Generally speaking, the dark blue nodes can be reached within the time threshold from base 6, and only base 6. Therefore, if the ambulance at base 6 is busy, incidents on the dark blue nodes will not be reached in time. For those dark blue nodes, the next closest base is base 3. But dispatching this vehicle (if it is idle) will leave the entire east side of the region without idle ambulances. Therefore, it is in this case better to use an ambulance from the west side of the region. The enlarged response time is -using our objective of the fraction late arrivals- not a downside, since the incident could not be reached in time anyway.

For incidents on the purple and cyan nodes, the best decision depends mostly on the state of the ambulance at base 3 and 6. If both ambulances are simultaneously busy, then the best ambulance to send to incidents in the purple region is usually the closest idle one. In the same scenario, incidents in the cyan region are typically *not* helped by the closest idle one. Note that this is the scenario when the entire east side of the region is not covered. This behaviour can be interpreted in a way similar to the case above (regarding dark blue nodes). When an incident cannot be reached in time, we might as well choose a vehicle other than the closest idle one. This can be beneficial, because the choice can be made such that the remaining ambulances are in a favourable position with respect to possible future incidents. Note that this is also the general idea that forms the basis of our MEXCLP dispatch heuristic.

7.3 Results

In this section, we show the results from our simulations of the EMS region of Utrecht.

We run simulations for region Utrecht, using four different dispatch policies: the closest idle method, the MEXCLP-based heuristic (using both destinations and real locations of vehicles) and the MDP solution after 21 value iterations. Figure 6 compares their performance in terms of the observed fraction of response times larger than the threshold time.

The results show that the MDP solution that was designed to minimize the fraction of late arrivals has

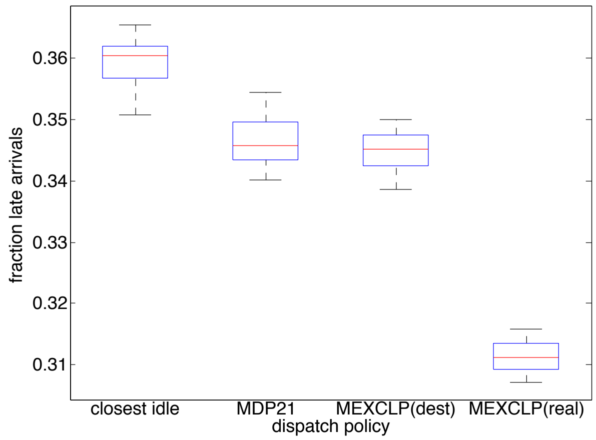

Figure 6: Comparing the performance of the MDP solution after 21 value iterations, with two variants of the Dynamic MEXCLP dispatch method (where $q = 0.3$). The benchmark is the 'closest idle' policy. Each policy was evaluated with 20 runs of 5,000 simulated hours.

approximately the same performance as the MEXCLP-based dispatch heuristic that uses the destinations of vehicles. Both policies perform better (on average) than the 'closest idle' policy. In addition, the MEXCLP-based dispatch heuristic that uses the real locations of vehicles performs even better.

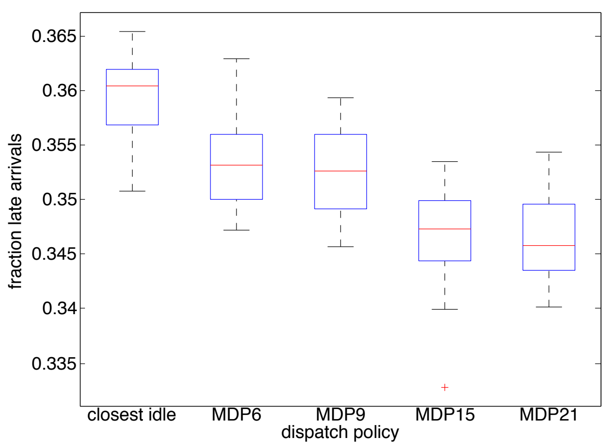

Figure 7: The performance of the MDP solution for region Utrecht after 6, 9, 15 and 21 value iterations. Each policy was evaluated with 20 runs of 5,000 simulated hours. The 'closest idle' dispatch policy is the benchmark.

For the region Utrecht with 8 ambulances, value iteration took a long time to converge. Instead of waiting for convergence, we applied the policy we get after a fixed number of value iterations. Figure 7 indicates that the performance increases when we increase the number of value iterations.

Up until now we have focused on the fraction of late arrivals, a key performance measure in ambulance operations. However, other aspects of the response times can also be important. For example, it is considered a drawback if patients have to wait an extremely long time for their ambulance to arrive (i.e., the response time distribution is heavy tailed). In this example -as well as in others- there exist trade offs between performance indicators.

We visualize the cumulative distribution of response times, as obtained from our simulation, in Figure 8. Figure 8 shows – just like Figure 6 – that the MEXCLP heuristic outperforms the other policies for response times within the time threshold (720 seconds). However, it also shows significant differences in response times for response times greater than T. for which the MEXCLP heuristic performs *worse* than the benchmark.

How much one is willing to sacrifice on one performance indicator in order to realize an improvement in the fraction of late arrivals, is typically the source of a lively discussion. Although such choices depend on how the different aspects of the response times are weighted, we expect that in realistic cases decision makers will prefer the closest idle policy over the MEXCLP heuristic. The reason for this is in the tail of the response times, see Figure 8.

When observing Figure 8, we make two other observations. First of all, the line of the MDP(frac)21 solution is very close to the MEXCLP(dest) line. Remember that these two policies base their decisions on the same information (that is, the destinations of idle vehicles and the location of a possible incident). This observation confirms our belief that these two policies have a similar underlying idea (they attempt to balance the response time for a current incident with the coverage for possible future incidents.) Secondly, one may note that the line for the MDP(avg) solution is remarkably similar to the line for the closest idle method. For this, we have no clear explanation (yet).

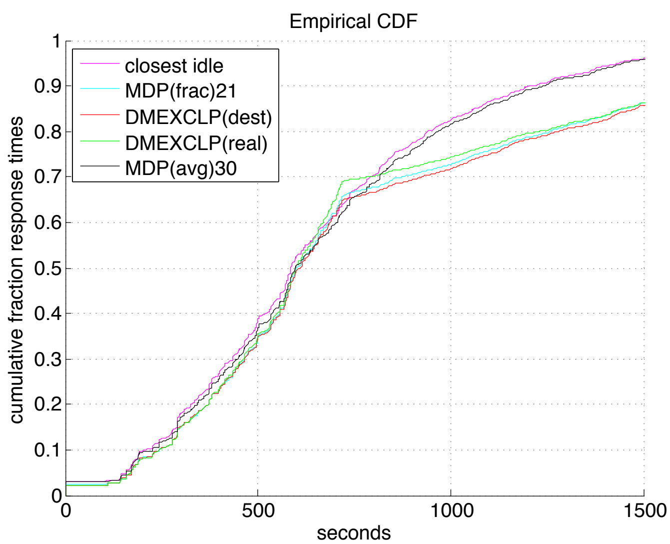

Figure 8: The cumulative distribution of response times observed in a simulation of 5,000 hours per dispatch policy, for the region Utrecht with 8 ambulances. For the MEXCLP algorithms, a value of $q = 0.2$ was used. For the Markov Decision Problems, the notation MDP(objective)#iterations was used.

7.4 Sensitivity to the parameter q

The dispatch heuristic based on MEXCLP has a parameter q, which represents the busy fraction. In this section we analyse the sensitivity of the performance to the value of q that is used. Thereto, we simulated the EMS system several times for several values of q.

In theory, q should be equal to the true busy fraction throughout the system. However, one may observe different behaviour for different values of q, and the true busy fraction need not necessarily be the one with optimal performance. This may seem counter-intuitive at first; but fact is that dynamic ambulance management is such a difficult problem, that we cannot hope to find a model that captures everything perfectly. Generally speaking, using MEXCLP with $q \to 0$ puts emphasis on covering the next incident. Using a higher busy fraction is equivalent with creating preparedness for incidents further into the future – at the cost of performing worse with respect to incidents in the near future. The true busy fraction could be a good starting point, but in practice one may choose a different value based on performance in simulations.

We simulated the EMS system of Utrecht, again with eight ambulances and (on average) four incidents per hour. We executed the MEXCLP dispatch heuristic for values of q between 0.1 and 0.8. The performance is shown in Figure 9. We observed that the true busy fraction throughout the simulations was between 37.5% and 38.1% (as measured when using $q = 0.2$ and $q = 0.8$ respectively).

Firstly, analysis of Figure 9 suggests that $q = 0.2$ would be a good choice for this particular scenario: it seems to result in the lowest fraction of late arrivals. Secondly, note that q varies between 0.1 and 0.8 in this analysis, which are quite extreme values. In practice, discussions will typically be about smaller perturbations, e.g., should we use $q = 0.2$ or $q = 0.3$? Furthermore, it is also important to recognize the scale on the vertical axis, as well as the overlap in the boxes of the box plot. Recall that the performance of the benchmark (the closest idle policy) is approximately 36%, which is significantly worse than our heuristic for any value of q. We conclude that the MEXCLP dispatch method is very insensitive to the value of the parameter q.

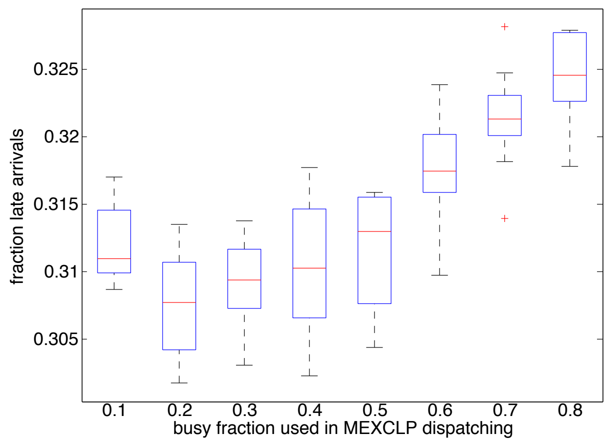

Figure 9: The performance observed for the MEXCLP dispatch heuristic, for different values of parameter q. Each box consists of 10 simulations of 5,000 hours each, for the region Utrecht with 8 ambulances.

8 Results: Region Utrecht with 19 vehicles

In the previous section, we used 8 vehicles in the region of Utrecht, due to the scaling limitations of our MDP solution. In this section, we analyze a more realistic representation of Utrecht: we increase the incident frequency to one incident per 10 minutes (on average). This is quite a reasonable estimate for this region during the summer period[2]. Simultaneously, we increase the total number of ambulances to 19. For the other simulation parameters, we use the same values as in Section 7.

We allow ambulances to be stationed only at locations that match the EMS base locations that exist in reality (using data from 2013). Throughout this section, we assign ambulances to the available bases according to the solution of the *static* MEXCLP model, which is generally assumed to give reasonable solutions (for a comparison of static methods, see (22)).

Figure 10 compares the performance of the MEXCLP dispatch heuristic with the benchmark (the closest idle policy). Note that the obtained fraction of late arrivals – roughly 5% – is realistic for this region in practice. The MEXCLP dispatch heuristic reduces the fraction of late arrivals from 0.053 to 0.043 (on average), a relative improvement of approximately 18%. To the best of our knowledge, no previous literature on ambulance dispatching has described a performance improvement of this magnitude – except perhaps for artificial problem instances that were designed for this purpose. Moreover, it was often assumed that changing the dispatch policy – as opposed to changing the position of idle vehicles – would not lead to major improvements (see, e.g., (23)). Our results shed new light on this belief. Note that an improvement of 18% is considered large, even with respect to algorithms that *are* allowed to reposition idle vehicles.

It should be clear that – when solely focusing on the fraction of late arrivals – the MEXCLP dispatch heuristic can offer great improvements compared to the closest idle policy. However, decision makers are often interested in more than just the fraction of late arrivals. They should be warned that changing from

[2]Our dataset for this region in the month of August 2008 shows 4775 urgent ambulance requests, which is on average 9.4 minutes between incidents.

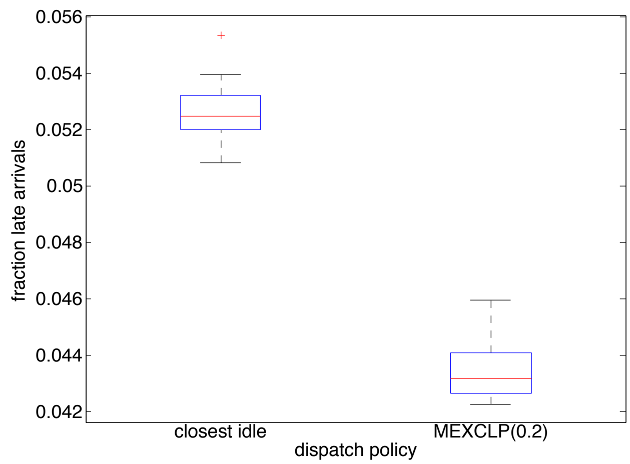

Figure 10: The objective (fraction of late arrivals), as observed in a simulation of 5000 hours per dispatch policy, for the region Utrecht with 19 ambulances.

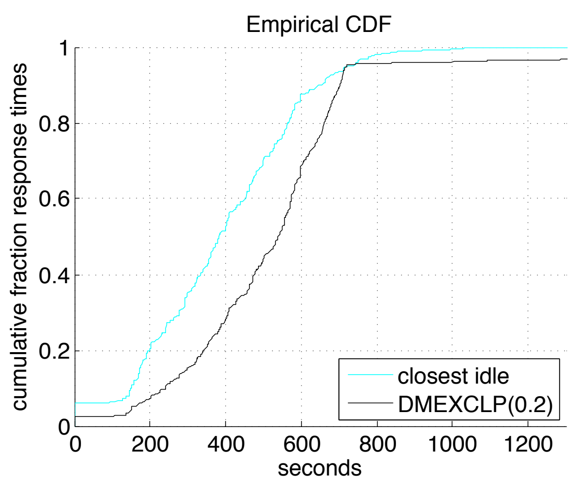

Figure 11: The cumulative distribution of response times, as observed in a simulation of 5000 hours per dispatch policy, for the region Utrecht with 19 ambulances.

the closest idle dispatch policy to the MEXCLP heuristic considerably diminishes the performance of other quality indicators, as can be seen in Figure 11. We highlight the difference in the average response time: when switching from the closest idle method to our heuristic, the average response time increased from 390 seconds to 535 seconds (an increase of 37%).

9 Conclusion

We have developed two methods to obtain ambulance dispatch policies.

Firstly, we modeled the ambulance dispatch problem as a Markov Decision Problem (MDP). This is the first MDP in ambulance literature that models more than just the number of idle vehicles, without losing tractability for reasonably-sized ambulance fleets. In some sense, this model balances the amount of detail in the system representation – which typically results in a better outcome – with the computational difficulties. The MDP can be used in two different settings: for minimizing the fraction of late arrivals or minimizing the average response time. The solution of the MDP performs slightly better than the benchmark (the closest idle policy).

Secondly, we approach the ambulance dispatch problem with an easy to implement heuristic. This heuristic reduces the expected fraction of late arrivals by 18% relatively compared to the benchmark, for a realistic EMS region. Our heuristic scales well: it could easily be applied to all realistic EMS regions, regardless of their size and fleet.

Limitations Although it is possible to apply our MDP in practice for reasonably-sized ambulance fleets, we do not recommend it: computation times are rather long and the performance improvement is small. The MDP is – in our opinion – mostly of theoretical interest. On the other hand, the heuristic could very well be applied in practice, but decision makers should be aware of its side effects: the heuristic aims to minimize the fraction of late arrivals, which does not reduce – and can in fact increase – response times overall.[3] We recognize this as an important downside and emphasize that practitioners should carefully consider whether the response time threshold really is the way they want to evaluate their performance.

Main contribution This paper sheds new light on the popular belief that deviating from the closest idle dispatch policy cannot greatly improve the objective (the expected fraction of late arrivals). We found an improvement of 18%, which was unexpectedly large. We consider this the main contribution of our work. Practitioners and researchers who define the fraction of late arrivals as their sole objective, should no longer claim that the closest idle method is near-optimal. Although we did not find dispatch policies that are practically preferable over the closest idle policy, we have shown that the argumentation for not using alternatives should be different. One should argue that they use the closest idle policy because we do not know alternatives that improve response times overall – and not because the alternatives fail to improve the fraction of late arrivals.

Acknowledgements

The authors of this paper would like to thank the Dutch Public Ministry of Health (RIVM) for giving access to the travel times for EMS vehicles in the Netherlands. This research was financed in part by Technology Foundation STW under contract 11986, which we gratefully acknowledge.

References

[1] R. Alanis, A. Ingolfsson, and B. Kolfal. A Markov chain model for an EMS system with repositioning. *Production and Operations Management*, 22(1):216–231, 2013.

[2] Damitha Bandara, Maria E. Mayorga, and Laura A. McLay. Optimal dispatching strategies for emergency vehicles to increase patient survivability. *International Journal of Operational Research*, 15(2):195–214, 8 2012.

[3] R. Bellman. *Dynamic Programming.* Princeton University Press, 1957.

[4] Richard Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, 6(4):679–684, 1957.

[5] R. Bjarnason, P. Tadepalli, and A. Fern. Simulation-based optimization of resource placement and emergency response. *Proceedings of the Twenty-First Innovative Applications of Artificial Intelligence Conference*, 2009.

[6] G. Carter, J. Chaiken, and E. Ignall. Response areas for two emergency units. *Operations Research*, 20(3):571–594, 1972.

[7] R.L. Church and C.S. Revelle. The maximal covering location problem. *Papers of the Regional Science Association*, 32:101–118, 1974.

[8] Personal communication with ambulance provider Utrecht (RAVU).

[3] Note that the same effect holds for the MDP that aims to minimize the fraction of late arrivals.

[9] M.S. Daskin. A maximum expected location model: Formulation, properties and heuristic solution. *Transportation Science*, 7:48–70, 1983.

[10] S. F. Dean. Why the closest ambulance cannot be dispatched in an urban emergency medical services system. *Prehospital and Disaster Medicine*, 23(02):161–165, 2008.

[11] J. Goldberg, R. Dietrich, J.M. Chen, and M.G. Mitwasi. Validating and applying a model for locating emergency medical services in Tucson, AZ. *Euro*, 34:308–324, 1990.

[12] C.J. Jagtenberg, S. Bhulai, and R.D. van der Mei. An efficient heuristic for real-time ambulance redeployment. *Operations Research for Health Care*, 4:27–35, 2015.

[13] J. P. Jarvis. Optimal assignments in a markovian queueing system. *Computers & Operations Research*, 8(1):17–23, 1981.

[14] R.B.O. Kerkkamp. Facility location models in emergency medical service: Robustness and approximations. Master's thesis, Delft University of Technology, 2014.

[15] Max Kuhn and Kjell Johnson. *Applied Predictive Modeling*. Springer New York, 2013.

[16] M.S. Maxwell, M. Restrepo, S.G. Henderson, and H. Topaloglu. Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing*, 22:226–281, 2010.

[17] Laura A. McLay and Maria E. Mayorga. A dispatching model for server-to-customer systems that balances efficiency and equity. *Manufacturing and Service Operations Management*, 15(2):205–220, 3 2013.

[18] Laura A. McLay and Maria E. Mayorga. A model for optimally dispatching ambulances to emergency calls with classification errors in patient priorities. *IIE Transactions (Institute of Industrial Engineers)*, 45(1):1–24, 1 2013.

[19] Ambulancezorg Nederland. Ambulances in-zicht 2014, 2014.

[20] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, NY, USA, 1994.

[21] C. Swoveland, D. Uyeno, I. Vertinsky, and R. Vickson. Ambulance location: A probabilistic enumeration approach. *Management Science*, 20(4):686–698, 1973.

[22] P.L. van den Berg, J.T. van Essen, and E.J. Harderwijk. Comparison of static ambulance location models. *Under review*, 2014.

[23] Yisong Yue, Lavanya Marla, and Ramayya Krishnan. An efficient simulation-based approach to ambulance fleet allocation and dynamic redeployment. In *AAAI Conference on Artificial Intelligence (AAAI)*, July 2012.